arXiv:1608.08080v1 [math.OC] 24 Aug 2016

A Multi-Urn Model for Network Search

Christopher E Marks
Operations Research Center, Massachusetts Institute of Technology
Charles Stark Draper Laboratory
555 Technology Square
Cambridge, MA 02139
cemarks@mit.edu

Tauhid Zaman
Department of Operations Management, Sloan School of Management
Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139
zlisto@mit.edu

We consider the problem of finding a specific target individual hiding in a social network. We propose a method for network vertex search that looks for the target vertex by sequentially examining the neighbors of a set of "known" vertices to which the target vertex may have connected. The objective is to find the target vertex as quickly as possible from amongst the neighbors of the known vertices. We model this type of search as successively drawing marbles from a set of urns, where each urn represents one of the known vertices and the marbles in each urn represent the respective vertex's neighbors. Using a dynamic programming approach, we analyze this model and show that there is always an optimal "block" policy, in which all of the neighbors of a known vertex are examined before moving on to another vertex. Surprisingly, this block policy result holds for arbitrary dependencies in the connection probabilities of the target vertex and known vertices. Furthermore, we give precise characterizations of the optimal block policy in two specific cases: (1) when the connections between the known vertices and the target vertex are independent, and (2) when the target vertex is connected to at most one known vertex. Finally, we provide some general monotonicity properties and discuss the relevance of our findings in social media and other applications.

1. Introduction

In this paper, we examine the problem of searching a social network for a particular target individual by sequentially examining the neighbors of other known users. Social media applications enable users to connect with each other, forming social networks. For different reasons which we will discuss, one may wish to find a target individual in the social network. One may have prior knowledge that the target individual is connected with a set of known users, and so the most logical place to begin searching is the neighbors of these known users. If querying each of these neighbors incurs some sort of cost, then the goal would be to find the target with as few queries as possible.

For example, suppose Mary is searching a social media application for an account belonging to John, an old friend from school with whom she has lost contact. From what she knows about John, Mary might be able to develop a list of accounts she knows about within the social media application to which John's account might be connected. For example, she might recall that John was good friends with Matt, who has a social media account that is known to Mary. She also might remember John was active in a certain charity, which also maintains a social media account known to Mary.

After developing such a list, Mary could sequentially explore each account's connections, but doing so could take a substantial amount of time. In order to find John's account quickly (assuming John has an account), Mary might devise a search strategy. For example, she might start by looking at accounts she feels are the most likely to be connected with John's account. Alternatively, she might start by looking at accounts with fewer connections, because her goal is to find John's account while minimizing the time spent exploring.

In this hypothetical scenario, what Mary is doing is an example of a network *vertex search* in which the sought object, or *target* might be found by examining the neighbors of a finite set of known vertices. Once the search target is found, the search typically terminates. Each known vertex i might have a different degree, N_i, requiring a different number of search queries to exhaustively search. From our scenario, we also consider that each known vertex i might have a different probability of being connected to the target vertex, which we denote as φ_i.

In this paper we present a probabilistic "multi-urn" model for searches of this nature in which we represent each known vertex as an urn containing a finite number of marbles. Each marble in an urn represents one of the respective vertex's neighbors. The search consists of successively drawing and and examining individual marbles from the urns with the goal of finding a red marble, representing the search target, in the fewest number of draws. Figure 1 depicts the multi-urn model for a network search on three known vertices. In this example, each of the known vertices is connected to the search target, so each urn contains a single red marble. Additionally, each urn contains a number of blue marbles that represent the other neighbors belonging to each respective vertex. Employing a dynamic programming framework, we provide some insight into the optimal search policy in this general model and under certain conditions.

1.1. Application

Our model applies to any scenario where one must sequentially search for the target amongst a set of entities which are separated into different clusters. In our vertex search problem, the entities are vertices and the clusters are neighborhoods of the known vertices. Our main motivation for this model is in social media applications where many times the goal is to find users who harass

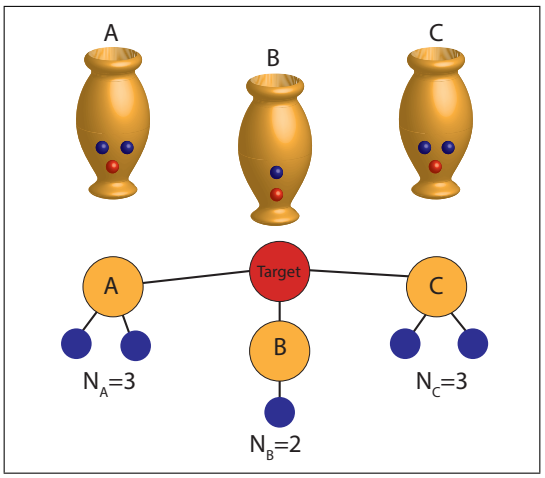

Figure 1 **Network search representation as a multi-urn model.**

others, incite violence, or engage in other dangerous behaviors. Twitter has been suspending large numbers of users, many of which support or engage in violent extremism, from its micro-blogging application for violating the site's published rules [11]. The challenge is that these users can simply create a new account each time one is suspended. However, from historical data Twitter could predict the accounts to which the suspended user is likely to connect. Using this information, Twitter administrators could then apply our optimality criteria to efficiently locate new accounts belonging to suspended users.

There are other scenarios where this model can apply. For instance, for law enforcement and intelligence applications, the search entities could be suspects in a crime and the clusters could be geographical locations. Or if one is examining a dump of emails from a suspect's server, one may be looking for an incriminating email, so the entities are emails and the clusters could be recipients of the emails. In both of these examples, the process of querying the entities requires a non-trivial amount of resources (interviewing a suspect, reading an email), so it is important to find the target as quickly as possible. Using an optimal or near-optimal search strategy is therefore crucial in these examples.

1.2. Our Contributions

In this paper we develop a multi-urn search model, a new and useful methodology for analyzing searches similar to the network vertex search problem we proposed in the introduction. We employ a dynamic programming framework to analyze this model and provide theoretical results based on the nature of the cost function.

In particular, we show that in this type of search problem there always exists an optimal policy, i.e., one that minimizes the expected number of marbles drawn before finding a red marble, in which the marbles in each urn are exhausted before moving on to the next urn (Theorem 2).

We refer to a policy that meets this criterion as a *block policy*, and show that this result holds irrespective of dependencies among the urns. This result is surprising because the presence of the red ball in the urns could have arbitrary correlations. There could be positive correlations, where if the ball is in one urn, it is more likely to be in another urn. Or the correlations could be negative, so if a ball is in one urn, it is less likely to be in another urn. Nonetheless, our result shows that no matter what the dependency between the urns, a block policy is always an optimal policy to find the red ball.

Building on this finding, we provide optimality conditions that enable immediate determination of an optimal policy in two specific cases:

- Each urn contains a red marble independent of other urns (Theorem 4). This case corresponds to the assumption that each known vertex is connected to the search target independent of other known vertices' connections.
- There is at most one red marble among all of the urns (Theorem 6). This case corresponds to the assumption that the search target is connected to at most one known vertex.

Finally, we provide detailed analysis of the system dynamics of the multi-urn search model, leading to insight into the intuition behind our findings. We provide monotonicity properties on the evolution of certain probabilities as marbles are successively drawn (Theorem 7) and establish a useful bound on how much the probability of drawing a red marble from a certain urn can change between two successive stages (Theorem 8).

1.3. Previous Work

Much of the work that has been done in the context of network search is focused on finding *relevant* vertices in a large scale network. Google's PageRank algorithm is probably the most well known example of these methods, of which many adaptations and generalizations exist [2]. Our work looks at an essentially different type of network search: one of finding a *specific* vertex in a network, presumably identifiable by certain features, by investigating network neighborhoods in which the vertex is likely to appear.

The network vertex search problem we have proposed could instead be formulated as a multi-arm bandit problem. Bubeck and Cesa-Bianchi [4] provide a broad survey of many variations of the multi-arm bandit problem and their respective applications. These problems are typically likened to a gambler who has a choice of playing from a set of slot machines. At each discrete stage in the process the gambler selects and plays a slot machine for a certain cost and receives a stochastic reward from an unknown distribution. The more times the gambler plays a particular machine, the more he is able to learn about its reward distribution.

In the multi-arm bandit setting, the gambler would not want to spend too much money playing low-payout slot machines just to learn their reward distribution. This quandary is the fundamental

tradeoff between *exploration* and *exploitation*, which is inherent in multi-arm bandit problems. In order to make money, the gambler wants to play only the highest-payout slot machine. However, he never really knows the true distributions of any of the machines. As a result, optimal multi-arm bandit policies often include a balance of exploratory actions, in which decisions are made for the sole purpose of observing outcomes, and exploitative actions, in which decisions are made to optimize the outcomes based on what has been learned.

The multi-arm bandit problem objective is often characterized as the minimization of *regret*, which is essentially the difference in expectation between what the gambler earns and what he would have earned by playing the best machine. Lai and Robbins [12] provide a very well-known method for constructing adaptive multi-arm bandit policies using upper confidence bounds, for which regret grows proportional to the logarithm of the number of plays in the limit. Auer et al. [1] show that this same bound on regret is also achievable in finite time.

The multi-urn search model we present could be cast in the context of a finite time multi-arm bandit problem, but there are a few notable differences. Our objective, to find the search target as quickly as possible, does not immediately cast itself as minimizing regret. Gittins [10] overcomes this difficulty by augmenting the state space in the multi-arm bandit formulation with a "success" state, from which no additional costs or rewards are incurred. Building on this adaptation, Gittins describes a class of search problems that are very similar to our network search problem, and characterizes the optimal search policy based on his well-known dynamic allocation index [9].

Our search problem differs from [10], however, in that each vertex has a fixed, finite, and known number of neighbors. In essence, we assume the reward distribution of each slot machine is known, and we only allow a fixed, finite number of plays on each machine. Unlike the bandit approach, the outcomes of successive marble draws from a single urn are not assumed to be independent observations from an unknown distribution. Instead, our model uses a known distribution on each machine but limits each machine to allowing at most a single win. Furthermore, in our approach we allow for dependencies between the urns, whereas the multi-arm bandit approach typically assumes each slot machine's outcomes are independent of the others.

In spite of these differences, the dynamic allocation index applied in the class of search models proposed by Gittins [10] has many similarities to our development. The system dynamics in both cases are governed by Bayesian probability updates. We show that in at least two cases the optimal policy can be characterized by a priority index, which is derived directly from the system dynamics and can be interpreted as the expected rewards of decisions. Gittins [10] also mentions monotonicity properties of his dynamic allocation index in the context of search that are similar to the monotonicity properties we derive. Our method for proving optimality uses similar logic to the proofs given in [8], which are based on the original development by Gittins and Jones [9].

Like multi-arm bandit problems, urn models have been applied in many contexts, including discrete decision processes. The Pólya urn process is a well-known construct using urns that has been adapted and used in many applications [14]. This process generally consists of one or more urns, each containing certain numbers of marbles of different colors. At each stage in the process a marble is randomly drawn from an urn and its color observed. This color then dictates an action involving placing one or more marbles of certain colors into certain urns.

Wei [15] provides a specific adaptation the Pólya urn process to the problem of conducting medical trials in a way that is meant to exploit the use of treatments that have shown positive results in the past, which is very similar to multi-arm bandit models applied in the same context. The Pólya urn process has also been used as a preferential attachment model in the formation of networks [6]. This application can be useful in considering how links form in social networks, and is similar to our problem. We assume, however, that the links are already present in the network and are instead interested in finding the optimal way to investigate these existing links.

Downey et al. [7] employ a multi-urn model that is very similar to ours but serves a different purpose: unsupervised information extraction. The model these authors propose uses urns to represent different collections of documents. Marbles drawn from the urns represent specific documents, from which specific labels are extracted. The objective of the model is to learn which labels are the correct, or "target" labels, and which labels are erroneous extractions.

The urn model proposed in [7] differs substantially from ours in its objective. Downey et al. have the objective of learning model parameters and, in the unsupervised case, learning which labels are correct. In the urn model we present, we assume the probability distributions and the target labels are known a priori, and we aim to to find a target marble as efficiently as possible.

Our network search problem is also related to the problem of mutual information maximization. If our goal was mutual information maximization, we would not necessarily focus our search effort on trying to find the target vertex. Instead, we would examine the places that would give us the most information about where the target is likely to be. This is similar to the goal of exploration in the multi-arm bandit problem. Chen et al. [5] analyze a sequential information maximization problem that parallels our development, using a dynamic programming approach and giving bounds on the performance of the greedy approach. The problem the authors propose involves learning about the distribution of an unknown parameter of interest by sequentially observing other variables. Each observation provides some information about the unknown parameter, and the objective is to maximize the total information gained in a fixed number of observations.

Our multi-urn search model departs most substantially from the development in [5] by imposing additional constraints and dynamics in the way observations are made. In our model, the urns are depleted over time, changing the amount of information contained in each successive marble drawn

in predictable, but sometimes unintuitive ways. Our main contributions are the characterizations of optimal search policies under various probability models, which come directly from analysis of the dynamics inherent in our multi-urn search model.

Finally, recent work in scheduling and inspection policies employ similar dynamic programming approaches to characterize optimal policies. Levi et al. [13] use dynamic programming to find policies that optimally allocate resources between information gathering and task execution. This class of models provides a natural extension to our network search problem. While we assume a probability model on a set of known vertices, using this approach we could attempt to find the optimal balance between the time spent learning a probability model on a set of known vertices and the time spent executing the search on the current known vertex set.

2. Multi-urn Search Model

We return to the context of network vertex search as presented in the introduction. Let $\mathcal{V}$ be the set of vertices to search, and let N_i be the number of search queries required to exhaustively search vertex $i \in \mathcal{V}$. We assume that the neighbors of each vertex are queried in a random order, so that each individual neighbor query of a particular vertex is equally likely to be the search target, given that the target is connected to the queried vertex.

Under these assumptions, we can represent this search problem as an experiment involving randomly drawing marbles from a set of urns, where each urn represents a known vertex in the network. Each marble in urn $i \in \mathcal{V}$ represents a neighbor of vertex i. The degree of vertex i is N_i, so urn i initially has N_i marbles. With probability $\varphi_i > 0$, exactly one of the N_i marbles in urn i is red, indicating that known vertex i is connected to the target vertex. Otherwise, all marbles in all urns are blue.

Querying a random neighbor of vertex i in search of the target is analogous to drawing a random marble from urn i and observing its color. If the marble is red, the target vertex has been located. If the marble is blue, the target has not been found and the search continues with the remaining marbles. Note that blue marbles are *not* put back into the urns; once they are drawn they are discarded. Just as Mary desires to find her old friend John with as few searches as possible, the goal in this experiment is to minimize the number of blue marbles drawn before finding a red one.

We now more completely specify the probability model that accounts for how the target vertex might be connected to the set of known vertices, i.e., how red marbles might be distributed among the urns. Let A_i be the event that the target vertex is connected to vertex i. We have already defined

$$\varphi_i = \mathbb{P}(A_i).$$

More generally, we let

$$\varphi_U = \mathbb{P}\left(\bigcap_{i\in U} A_i\right)$$

be the probability that the target vertex is connected to all vertices in set $U \subseteq \mathcal{V}$. If we were to assume that the target would connect to the members of U independently, then $\varphi_U = \prod_{i\in U} \varphi_i$. In general, the connections might not be independent. For example, Mary might think that if John connected with a certain musician he liked, he might be more likely to connect to other, similar musicians. In other cases, a connection to a particular vertex might imply a decrease in the probability of connection to another vertex.

In our urn model, we assume a known probability φ_U for all subsets $\{U : U \subseteq \mathcal{V}\}$, which fully specifies a probability model on the locations of the red marbles among the urns. It allows for arbitrary correlations between urns, so that the presence of a red marble in one urn (or subset of urns) can have a positive or negative correlation with the presence of a red marble in another urn (or another subset of urns).

We note now that the empty set $\emptyset \in \{U : U \subseteq \mathcal{V}\}$. By convention, we set $\bigcap_{i\in\emptyset} A_i = \Omega$, so that $\varphi_\emptyset = 1$. This term is implicitly included in summations over all subsets expressed in this paper. For example, the summation

$$\sum_{U\subseteq\mathcal{V}} (-1)^{|U|}\varphi_U$$

includes a "1" corresponding to the case in which $U = \emptyset$.

Given this set of probabilities, the probability of any specific outcome of marble locations, or vertex connections, can be determined using the well-known inclusion-exclusion formula. For example, suppose we are interested in the probability that the marble is located in all of the urns in set U and no other urns. This event can be written as $\left(\bigcap_{i\in U} A_i\right) \cap \left(\bigcap_{j\in\mathcal{V}\setminus U} A_j^c\right)$, with

$$\mathbb{P}\left(\left(\bigcap_{i\in U} A_i\right) \cap \left(\bigcap_{j\in\mathcal{V}\setminus U} A_j^c\right)\right) = \sum_{S:S\subseteq\mathcal{V},\ U\subseteq S} (-1)^{|S|-|U|}\varphi_S \geq 0. \tag{1}$$

Throughout this paper, we refer to the type of search described in this section as a *multi-urn search problem* which we now more formally define.

Definition 1 *A* multi-urn search problem *is a search problem that can be modeled as sequentially drawing marbles from a set of urns, $\mathcal{V}$, where*

1. *The objective of the searcher is to find a red marble with as few draws as possible.*
2. *Each urn $i \in \mathcal{V}$ contains at most a single red marble. Otherwise, all marbles are blue.*
3. *Each urn $i \in \mathcal{V}$ contains a known number of marbles, N_i.*

4. *For each subset of urns $U \subseteq \mathcal{V}$, the probability that a red marble is present in all urns in U is φ_U. Additionally, we set $\varphi_\emptyset = 1$.*

The network vertex search problem we used to motivate this model can be characterized as a multi-urn search problem, but this model might have other useful applications as well. For this reason, in the remainder of this paper we provide all analyses and results in the multi-urn search context, using the language of "urns" and "marbles," though we could immediately recover our original context by substituting "known vertices" and "neighbors," respectively.

2.1. Dynamic Programming Framework

In the search model we have defined, the decisions are carried out sequentially in discrete stages. We now take a dynamic programming approach [3] to framing this problem.

We model the search process as a discrete dynamic system of the form

$$\mathbf{x}(t+1) = f(\mathbf{x}(t), u(t), w(\mathbf{x}(t), u(t))),$$

where $t = 0, 1 \ldots$ is the *stage* of the search, which we equate to the total number of marbles already drawn from the urns. The system state, $\mathbf{x}(t)$, is a record of the total number of marbles drawn from each urn, which sufficiently characterizes the system at stage t. Parameter $u(t)$ is the decision made, or urn selected, at stage t, and $w(\mathbf{x}(t), u(t))$ is a binary stochastic input that indicates whether a red marble is drawn from urn $u(t) \in \mathcal{V}$ in state $\mathbf{x}(t)$.

If a red marble is drawn at stage t, then $w(\mathbf{x}(t), u(t)) = 1$ and the search terminates. Otherwise, $w(\mathbf{x}(t), u(t)) = 0$ and the search continues. Letting $x_i(t)$ be the number of marbles that have been removed from urn i at time t, we can explicitly define the state vector

$$\mathbf{x}(t) = (x_1(t), x_2(t), \ldots, x_{|\mathcal{V}|}).$$

If a blue marble is drawn from urn $u(t)$ in state $\mathbf{x}(t)$, the state transition function is:

$$f(\mathbf{x}(t), u(t), 0) = \mathbf{x}(t) + \mathbf{e}_{u(t)},$$

where $\mathbf{e}_i$ is the ith unit vector. If a red marble is drawn at any stage, the search terminates.

Our dynamic programming model consists of at most $N+1$ stages, where $N = \sum_{i \in \mathcal{V}} N_i$ is the total number of marbles summed over all of the urns, and is finite.

We define a valid *policy* $\mathbf{u}$ as a sequence of decisions $(u(0), u(1), \ldots, u(N-1))$, where $u(t) \in \mathcal{V}$ for $t = 0, 1, \ldots, N-1$, *and* for which

$$|\{t : u(t) = i\}| = N_i \qquad \forall\ i \in \mathcal{V}.$$

This final condition ensures that the policy will eventually exhaust each urn, as long as a red marble is not found, while at the same time never attempting to draw marbles from an empty urn. A searcher executing a valid policy draws a marble from urn $u(t)$ at each stage t until either the target marble is found or there are no marbles remaining in any of the urns, in which case the entire policy has been executed.

We note that in this dynamic programming model there is no benefit in making policy decisions during search execution. At each stage the searcher either draws a red marble, in which case she stops looking, or draws a blue marble and keeps searching. A valid policy provides an ordering of urn queries that is essentially conditioned on not drawing a red marble, which can be considered a deterministic process governed by our simple state transition function. The expected search outcomes for such a policy can be analyzed and compared to those of other policies a priori.

2.1.1. Stage t Probability of Drawing a Red Marble Building on our dynamic programming modeling assumptions, we now develop the probability distribution associated with $w(\mathbf{x}(t), u(t))$. Recall that this function indicates whether a red marble is drawn in stage t: $w(\mathbf{x}(t), u(t)) = 1$ implies a red marble is drawn from urn $u(t)$ at stage t, while $w(\mathbf{x}(t), u(t)) = 0$ implies a blue marble is drawn from $u(t)$ at stage t.

In determining the probability distribution of $w(\mathbf{x}(t), u(t))$, it is important to remember that in order to arrive in stage t while executing policy $\mathbf{u}$, the preceding queries $u(0), u(1), \ldots, u(t-1)$ would have been executed *without drawing a red marble*, so that the system arrives in state $\mathbf{x}(t)$. For simplicity of notation, we condition an event on state $\mathbf{x}(t)$ to imply that state $\mathbf{x}(t)$ has been reached without having drawn a red marble. For example, $\mathbb{P}(A_i \mid \mathbf{x}(t))$ represents the probability urn i contains a red marble, given queries $u(0), u(1), \ldots, u(t-1)$ have been executed without drawing a red marble.

Using the multiplication rule, we can write the probability

$$\mathbb{P}(w(\mathbf{x}(t), u(t)) = 1) = \left(\frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}(A_{u(t)} \mid \mathbf{x}(t)), \tag{2}$$

which is the probability of drawing a red marble from the $N_{u(t)} - x_{u(t)}(t)$ marbles remaining in urn $u(t)$, given there is a red marble in $u(t)$, multiplied by the probability urn $u(t)$ contains a red marble given the system has arrived at state $\mathbf{x}(t)$.

The complementary probability can be written using the law of total probability:

$$\begin{aligned}\mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) &= \left(1 - \frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}(A_{u(t)} \mid \mathbf{x}(t)) + \mathbb{P}(A^c_{u(t)} \mid \mathbf{x}(t)) \\ &= 1 - \left(\frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}(A_{u(t)} \mid \mathbf{x}(t))\end{aligned} \tag{3}$$

2.1.2. Stage t Urn Probabilities In this process, we have assumed a fully specified initial probability model on the urns, i.e., for any subset $U \subseteq \mathcal{V}$, the probability that a red marble is present in all of the urns, φ_U, is known. This probability model can be thought of as a Bayesian prior, a quantification of the searcher's beliefs on where a red marble might be found.

However, these probabilities are not static. After drawing a marble from an urn, the probabilities change as a result of the new information. If the marble drawn is red, then the probability that a red marble existed in the queried urn becomes 1. Likewise, if the marble drawn is blue, then the probability that a red marble can be found in the queried urn decreases as a function of the number of marbles remaining in the urn and the current urn probability.

As long as a red marble is not found, the evolution of urn probabilities over the course of the search is completely determined by the initial probability model and the search policy. We now provide a general expression for updated urn probabilities at stage t.

Theorem 1 (Urn Probabilities). *In a multi-urn search problem over a set of urns $\mathcal{V}$, suppose a red marble is not found in the first t queries when executing a valid policy $\mathbf{u} = (u(0), \ldots, u(N-1))$. Then, for any subset of urns $U \subseteq \mathcal{V}$, the probability of a red marble being in all of the urns in U at stage t is given by:*

$$\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t)\right) = \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}.$$

Proof. See Section 4.1.

Substituting the result in Theorem 1 into Equations (2) and (3) gives us the following corollary.

Corollary 1 *The probability distribution of $w(\mathbf{x}(t), u(t))$, indicating whether a red marble is drawn at stage t is,*

$$\mathbb{P}(w(\mathbf{x}(t),u(t)) = k) = \begin{cases} \dfrac{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{i\in S}\frac{x_i(t+1)}{N_i}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_i(t)}{N_i}}, & k=0 \\ \left(\frac{1}{N_{u(t)}}\right)\dfrac{\sum_{\{S\subseteq\mathcal{V}:\, u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{i\in S\setminus\{u(t)\}}\frac{x_i(t)}{N_i}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{i\in S}\frac{x_i(t)}{N_i}}, & k=1. \end{cases}$$

Theorem 1 provides a few important insights. First, we can see that the urn probabilities at any stage depend only on the numbers of marbles drawn from all of the urns, and not the order in which they were drawn. This *path-independence* property of the stage t urn probabilities is somewhat intuitive: a specific set of marbles drawn gives us a fixed amount of information irrespective of the order in which we inspect the marbles. We will make use of the path-independence property in the proofs for Theorems 2, 4, and 6.

Another observation is that the form of the probability expression is similar to the well-known inclusion-exclusion formula for computing probabilities of unions of events. In fact, this probability is an application of the principle of inclusion-exclusion applied in conjunction with Bayesian updates. In Lemma 1, we explicitly define the events that are characterized by the inclusion-exclusion formulas in Theorem 1.

2.2. Costs

In many network search applications, the cost of finding and examining a (random) neighbor of a known vertex is primarily the time consumed in executing the query and reviewing the results to determine whether the neighbor is the search target. Because we have no reason to believe this time-cost would be different for different vertex-neighbor queries, we assume in our model that the cost of drawing a marble is the same for all urns. The goal of the searcher is simply to minimize the number of blue marbles drawn, or vertex-neighbor queries executed, before finding the search target.

We therefore define the cost function at stage t,

$$g(t) = \begin{cases} 1 & w(\mathbf{x}(t), u(t)) = 0 \\ 0 & \text{otherwise,} \end{cases}$$

which applies a unit cost for every blue marble drawn. Because this quantity is stochastic, we set as our objective the minimization of expected total cost. Letting random variable $C = \sum_{t=0}^{N} g(t)$, we aim to find the optimal policy $\mathbf{u}$ to solve the following optimization problem:

$$\text{minimize}_{\mathbf{u}} \ \mathbb{E}[C].$$

Because $C \in \{0, 1, \ldots, N\}$ almost surely, we can write

$$\begin{aligned} \mathbb{E}[C] &= \sum_{k=0}^{N-1} \mathbb{P}(C > k) \\ &= \sum_{k=0}^{N-1} \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0). \end{aligned} \tag{4}$$

The product in this summation, $\prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0)$, is exactly the probability of making it to stage $k+1$ without having found a red marble. From Corollary 1 we can find an expression for this probability.

Corollary 2 *Given a multi-urn search problem and valid policy* $\mathbf{u}$*, the probability of arriving in stage* $k+1$ *without having found a red marble is*

$$\begin{aligned} \mathbb{P}(C > k) &= \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) \\ &= \sum_{S \subseteq \mathcal{V}} (-1)^{|S|} \varphi_S \prod_{i \in S} \frac{x_i(k+1)}{N_i}. \end{aligned}$$

We can therefore rewrite the cost function,

$$\mathbb{E}[C] = \sum_{t=0}^{N-1} \sum_{S \subseteq \mathcal{V}} (-1)^{|S|} \varphi_S \prod_{i \in S} \frac{x_i(t+1)}{N_i}. \tag{5}$$

Substituting the probability from Corollary 2 into Corollary 1 reveals an interesting property of the dynamics of this system:

$$\mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) = \frac{\mathbb{P}(C > t)}{\mathbb{P}(C > t-1)}. \tag{6}$$

3. Key Results

The cost function in equation (5) is nonlinear and nonconvex. Additionally, for a solution to be feasible, the values for $x_i(t)$, $i = 1, \ldots, |\mathcal{V}|$, $t = 0, \ldots, N$ must be constrained to correspond to stages reached by a valid policy. Nonlinear, non-convex constrained optimization is difficult in general. However, the structure of the cost function enables us to derive some useful results that characterize the optimal solution in general, and provide necessary and sufficient conditions for optimality in some specific cases.

3.1. Block Policy Optimality

We now provide our primary general result, in which we give a characterization of an optimal search policy in the multi-urn search problem. We begin with a definition.

Definition 2 *A* block policy *is a valid policy* $\mathbf{u}_B$ *in which each urn is queried exhaustively prior to querying another urn. A block policy can be specified as a sequence of urns* $\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$, $v^i \in \mathcal{V}$, *implying*

$$u(t) = v^i, \quad \sum_{j=1}^{i-1} N_j \leq t < \sum_{j=1}^{i} N_j.$$

This definition can be used to characterize the optimal policy, which we now formally state.

Theorem 2 (Block Policy Optimality). *Given a multi-urn search problem in which the objective is to minimize the expected number of searches required to find a red marble, an optimal search policy exists that is a block policy.*

Proof. See Section 4.2.

This result says that an optimal policy for the multi-urn search problem can be characterized by a sequence of urns. Once this is specified, one then simply searches each urn until it is out of balls or a red ball is found. The surprising part of this result is that this block policy optimality holds for arbitrary correlations in the a priori connection probabilities. For instance, there can be

a negative correlation between two urns, where if the red ball is more likely to be in one urn, it is less likely to be in another. In this case one may intuitively expect that after querying an urn many times and not finding a red ball, at some point it might be advantageous to search another urn which has a negative correlation with the queried urn. However, our result says that it is optimal to continue querying the current urn until it is exhausted.

While Theorem 2 allows for optimal policies that are not block policies, constructing such a case requires initial conditions that include probabilities that are zero. If $\varphi_{\{i,j\}} > 0$ for all pairs $\{i,j\} \subset \mathcal{V}$ (as in the case of independent urns), then only block policies can be optimal policies. This result follows from the proof of Theorem 2 (Section 4.2): observe that this condition implies that function $h(t)$ in equation (13) is strictly increasing in t, creating a contradiction in equation (14).

We have shown that for mutli-urn search problems, a block policy is optimal, but we have not yet specified what the block policy is. In general it can be difficult under arbitrary correlation structures to find the optimal policy. However, under certain assumptions on the urn probability model, explicit necessary and sufficient optimality conditions can be found. We examine these conditions next.

3.2. Independent Urns

We now consider the special case in which the red marbles are assumed to be independently present in each of the urns, so that the presence of a red marble in any urn (or group of urns) does not affect the probability of a red marble being present in any other urn (or group of urns). This probabilistic independence can be formalized mathematically.

Definition 3 *An* independent *multi-urn search problem is a multi-urn search problem in which, for any subset of urns, $U \subseteq \mathcal{V}$,*

$$\varphi_U = \prod_{i \in U} \varphi_i.$$

Intuitively, this independence property should be maintained throughout the search process for any search policy, as we now show explicitly.

Theorem 3 (Independent Urn Probabilities). *Given an independent multi-urn search problem, then for any policy* $\mathbf{u}$*, at any stage t, the independence property is maintained so that*

$$\mathbb{P}\left(\bigcap_{i \in U} A_i \,\middle|\, \mathbf{x}(t) \right) = \prod_{i \in U} \mathbb{P}(A_i | \mathbf{x}(t)).$$

Proof. See Section 4.3

It follows from the result in Theorem 3 that the probability of finding a red marble in stage t is only a function of the initial conditions and number of times $u(t)$ has been queried in the past. The number of marbles that have previously been drawn from other urns $i \neq u(t)$ do not affect $\mathbb{P}(w(\mathbf{x}(t), u(t)) = 1)$.

Because of the independence of the urn probabilities, we are able to obtain closed form expressions for the probability of finding a red ball and the expected cost of a block policy, which are stated in the following results.

Corollary 3 *Given an independent multi-urn search problem, the probability distribution of* $w(\mathbf{x}(t), u(t))$ *at any stage* t *is*

$$\mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) = \frac{N_{u(t)} - x_{u(t)}(t+1)\varphi_{u(t)}}{N_{u(t)} - x_{u(t)}(t)\varphi_{u(t)}}.$$

Corollary 4 *Given an independent multi-urn search problem and a block policy* $\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$, $v^i \in \mathcal{V}$, *such that* $\tau(i) = \sum_{j=1}^{i-1} N_{v^j}$ *is the first stage in which urn* v^i *is queried. Then,*

$$\prod_{t=\tau(i)}^{\tau(i)+N_i-1} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) = (1 - \varphi_{v^i}),$$

the contribution of urn v^i *to the total expected cost is*

$$\sum_{k=\tau(i)}^{\tau(i)+N_i-1} \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) = \left(N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2} \right) \prod_{j=1}^{i-1} (1 - \varphi_{v^j}),$$

and the total expected cost is

$$\mathbb{E}[C] = \sum_{i=1}^{|\mathcal{V}|} \left(N_i - \frac{(N_i+1)\varphi_i}{2} \right) \prod_{j=1}^{i-1} (1 - \varphi_j).$$

Independence implies that knowing the composition of marbles in urn i does not provide any additional information on the compositions of marbles in any of the other urns. Drawing a marble from urn $u(t)$ in stage t still results in an update to this urn's probability in stage $t+1$, but all other urn probabilities remain stationary in this state transition. This property enables us to characterize the optimal policy in the case of independent urns.

Theorem 4 (Independent Urns Optimality) *Given an independent multi-urn search problem, a block policy*

$$\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$$

is optimal if and only if

$$N_{v^i} \left(\frac{2 - \varphi_{v^i}}{\varphi_{v^i}} \right) \leq N_{v^{i+1}} \left(\frac{2 - \varphi_{v^{i+1}}}{\varphi_{v^{i+1}}} \right), \quad i = 1, 2, \ldots, |\mathcal{V}|. \tag{7}$$

Proof. See Section 4.4.

We note that this policy is not greedy, i.e., it does not maximize the probability of finding the red marble at each stage. Rather, the optimality condition in equation (7) balances the probability of immediately drawing a red marble with the probability of finding a red marble in successive draws from the same urn.

To gain intuition, consider a two-urn example in which each urn has the same probability of containing a red marble ($\varphi_1 = \varphi_2$), but urn 1 has fewer marbles ($N_1 < N_2$). In this case the optimality condition in equation (7) would have us initially draw marbles from urn 1, which has the same probability of giving us the red marble as urn 2 but requires fewer draws.

Alternatively, consider the case in which the two urns have the same number of marbles but $\varphi_1 < \varphi_2$. In this case, the optimal policy according to equation (7) would have us draw from urn 2 first, which is more likely than urn 1 to give us a red marble in the same number of draws.

In order to more clearly distinguish between a greedy policy and an optimal one, we provide one more example. Consider the following independent multi-urn search problem with two urns. Urn 1 contains $N_1 = 1$ marble and has probability $\varphi_1 = \frac{9}{16}$ of containing a red marble. Urn 2 contains $N_2 = 2$ marbles and has probability $\varphi_2 = 1$ of containing a red marble. This problem admits two block policies: $\mathbf{u}_B = (2, 1)$ and $\tilde{\mathbf{u}}_B = (1, 2)$.

Policy $\tilde{\mathbf{u}}_B$ is a greedy policy; urn 1, which has the highest immediate probability of producing a red marble ($\frac{9}{16}$), is queried before urn 2. The expected number of blue marbles drawn using policy $\tilde{\mathbf{u}}_B$ is

$$\mathbb{E}[\tilde{C}] = \frac{7}{16} + \left(\frac{7}{16}\right)\left(\frac{1}{2}\right) = \frac{21}{32}.$$

Alternatively, if we follow policy $\mathbf{u}_B$ and draw from urn 2 first, then expected cost is

$$\mathbb{E}[C] = \frac{1}{2},$$

which is optimal. By accepting a slightly lower probability in the first draw, this policy guarantees that the red marble is found in at most two draws. The optimality condition in Theorem 4, equation (7) provides the best balance between the immediate and long-term benefits of each query.

3.3. One Marble

We now turn our attention to another special case of the multi-urn search problem in which we limit the total number of red marbles among all of the urns to a single marble. In our network search scenario, this constraint would follow from assuming that the target user is connected to *at most* one of the known accounts on Mary's list. This might not be a valid assumption for Mary to make, but it might apply to other search scenarios both in and out of the network context.

For example, suppose law enforcement investigators have evidence that a suspect made a single phone call from an unknown phone number during a certain period. Having obtained phone records from all likely recipients, they must efficiently search for the phone call of interest within the records of these likely recipients.

For a non-network example, suppose a hotel custodian, after servicing all of the hotel rooms, realizes he left his car keys in one of the rooms. The hotel might consist of several wings, each with different numbers of rooms, and the custodian might feel the loss was more probable in certain wings. The custodian wants to search the rooms efficiently for his keys, in order to find them before new customers begin to arrive.

This one-marble constraint imposes the strongest negative correlations between the urns: if the red marble is in urn i then it cannot be in j, i.e.,

$$P(A_j|A_i) = 0, \qquad i \neq j.$$

Another way to characterize this constraint is to state that the events $A_1, A_2, \ldots, A_{|\mathcal{V}|}$ are disjoint. We now formalize this notion in a definition.

Definition 4 *A single marble multi-urn search problem is a multi-urn search problem for which*

$$\varphi_U = 0 \qquad \forall\ U \subseteq \mathcal{V} \text{ such that } |U| > 1.$$

We now analyze of the single marble search problem. First we observe that

$$\sum_{i \in \mathcal{V}} \varphi_i \leq 1.$$

We allow for the possibility that this sum is strictly less than one, implying there is a chance that none of the urns contain the red marble. If the sum is equal to one, then the assumption is that exactly one of the urns contains one red marble. Theorem 5 specifies the single marble urn probabilities for an arbitrary state $\mathbf{x}(t)$.

Theorem 5 (Single Marble Urn Probabilities). *Given a single marble multi-urn search problem and a search policy* $\mathbf{u}$*. Then, the probability that a red marble is in urn* i *given state* $\mathbf{x}(t)$*, and given no red marble has been found in the first* t *queries, is*

$$\mathbb{P}(A_i \,|\, \mathbf{x}(t)) = \frac{\left(1 - \frac{x_i(t)}{N_i}\right)\varphi_i}{1 - \sum_{j \in \mathcal{V}} \frac{x_j(t)\varphi_j}{N_j}}.$$

The probability that the red marble is in all of the urns in any subset $U \subset \mathcal{V}$*, where* $|U| > 1$ *is*

$$\mathbb{P}\left(\bigcap_{i \in U} A_i \,\middle|\, \mathbf{x}(t)\right) = 0.$$

Proof. This result follows immediately from the definition of a single marble multi-urn search problem and Theorem 1. □

Because the red ball can only be in one urn, all joint probabilities are zero. This greatly simplifies our analysis and allows us to obtain closed form expressions for the probability of finding a red ball and the expected cost of a block policy, which are stated in the following results.

Corollary 5 *Given a single marble multi-urn search problem and a valid search policy* $\mathbf{u}$*, the probability distribution of* $w(\mathbf{x}(t), u(t))$*, conditioned on not having found a red marble in a previous stages, is*

$$\mathbb{P}(w(\mathbf{x}(t),u(t))=k)=\begin{cases}\frac{1-\sum_{j\in\mathcal{V}}\left(\frac{x_j(t+1)\varphi_j}{N_j}\right)}{1-\sum_{j\in\mathcal{V}}\left(\frac{x_j(t)\varphi_j}{N_j}\right)}, & k=0\\ \frac{1}{N_{u(t)}}\left(\frac{\varphi_{u(t)}}{1-\sum_{j\in\mathcal{V}}\left(\frac{x_j(t)\varphi_j}{N_j}\right)}\right), & k=1.\end{cases}$$

Corollary 6 *Given a single marble multi-urn problem and a block policy* $\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$*,* $v^i \in \mathcal{V}$*, such that* $\tau(i) = \sum_{j=1}^{i-1} N_{v^j}$ *is the first stage in which urn* v^i *is queried. Then, the probability of not finding the red marble before reaching stage* $\tau(i)$ *is*

$$\prod_{t=0}^{\tau(i)}\mathbb{P}(w(\mathbf{x}(t),u(t))=0)=1-\sum_{j=1}^{i-1}\varphi_{v^j},$$

the contribution of urn v^i *to the total expected cost is*

$$\sum_{k=\tau(i)}^{\tau(i)+N_i-1}\prod_{t=0}^{k}\mathbb{P}(w(\mathbf{x}(t),u(t))=0)=\left(N_i-\frac{(N_i+1)\varphi_i}{2}-N_i\sum_{j=1}^{i-1}\varphi_j\right),$$

and the total expected cost is

$$\mathbb{E}[C]=\sum_{i=1}^{|\mathcal{V}|}\left(N_{v^i}-\frac{(N_{v^i}+1)\varphi_{v^i}}{2}\right)-\sum_{i=1}^{|\mathcal{V}|-1}\sum_{j=i+1}^{|\mathcal{V}|}N_{v^j}\varphi_{v^i}.$$

Theorem 6 characterizes the optimal solution in the single marble case.

Theorem 6 (Single Marble Optimality). *Given a single marble multi-urn search problem, a block policy*

$$\mathbf{u}_B=(v^1,v^2,\ldots,v^{|\mathcal{V}|})$$

is an optimal policy if and only if

$$\frac{\varphi_{v^i}}{N_{v^i}}\geq\frac{\varphi_{v^{i+1}}}{N_{v^{i+1}}},\quad i=1,2,\ldots,|\mathcal{V}|. \tag{8}$$

Proof. See Section 4.5.

The optimality condition given in equation (8) leads to a greedy policy in which, at each stage, the marble that is drawn is the one that is most likely to be red. At stage $t=0$ this is certainly true, as the probability of drawing a red marble from any urn i in the first draw is φ_i/N_i, which is exactly what the optimality condition optimizes. In the next section we show that that this condition is maintained through state transitions in an optimal policy. If drawing a marble from urn i has the highest probability of producing a red marble in stage t, then (assuming the urn has at least one marble remaining) drawing another marble from the same urn maximizes the probability of finding a red marble in stage $t+1$.

3.4. Monotonicity Properties

We now provide a few monotonicity properties that give additional insight into the dynamics of multi-urn search problems, as well as block policy optimality.

Theorem 7 (Monotonicity). *Given a multi-urn search problem and a search policy* $\mathbf{u}$*, the following inequalities hold:*

1. *For any subset of urns* $U \subseteq \mathcal{V}$ *such that* $u(t) \in U$*,*

$$\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t)\right) \geq \mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t+1)\right),$$

with equality holding only in cases in which $\mathbb{P}\left(A_{u(t)} \mid \mathbf{x}(t)\right) = 1$ *or* $\mathbb{P}\left(\bigcap_{i\in U} A_i \mid \mathbf{x}(t)\right) = 0$*.*

2. *For any stage* t *for which urn* $u(t)$ *has more than one marble remaining, i.e.,* $N_{u(t)} - x_{u(t)}(t) > 1$*,*

$$\mathbb{P}(w(\mathbf{x}(t), u(t)) = 1) \leq \mathbb{P}(w(\mathbf{x}(t+1), u(t)) = 1),$$

with equality holding only when $\mathbb{P}(w(\mathbf{x}(t), u(t)) = 1) = 0$*.*

Proof. See Section 4.6.

These monotonicity properties provide intuition into why optimal block policies exist. Suppose at stage t a blue marble is drawn from urn $u(t)$. At stage $t+1$, the probability of urn $u(t)$ containing a red marble has decreased as a result of this new information. However, the probability that the next marble drawn from urn $u(t)$ is red has *increased* from the previous stage. If drawing from urn $u(t)$ in stage t had a high probability of returning a red marble, drawing another marble from $u(t)$ in stage $t+1$ has an even higher probability of producing a red marble.

In the case of independent urns, this property provides justification for using a block policy. Suppose urn $u(t)$ has multiple marbles in it and is optimal at stage t, and a blue marble is drawn from this urn. At stage $t+1$ the probability of drawing a red marble from urn $u(t)$ has increased, while all other urn and marble probabilities have remained unchanged from the previous stage t. It follows that it would continue to be optimal to draw from urn $u(t)$.

If we allow for correlations among the urns, however, drawing a blue marble from urn $u(t)$ might also increase the probability of drawing a red marble from other urns in the following stage. In the single marble multi-urn search problem, drawing a blue marble from urn $u(t)$ in stage t increases the probability of finding a red marble in each of the other urns in stage $t+1$. We now provide our final theoretical result, which states that the rate of probability growth in a queried urn is always at least as large as the rate of probability growth in any other urn.

Theorem 8 (Marble Probability Bound). *Given a multi-urn search problem and a search policy* $\mathbf{u}$*, the following inequality holds:*

$$\frac{\mathbb{P}\left(w(\mathbf{x}(t+1),u(t))=1\right)}{\mathbb{P}\left(w(\mathbf{x}(t),u(t))=1\right)} \geq \frac{\mathbb{P}\left(w(\mathbf{x}(t+1),i)=1\right)}{\mathbb{P}\left(w(\mathbf{x}(t),i)=1\right)}.$$

Proof. See Section 4.7.

Theorem 8 provides much intuition about why optimal block policies always exist in multi-urn search problems. It also shows that a purely greedy strategy, in which the probability of immediately drawing a red marble is maximized at each stage, will always produce a (possibly suboptimal) block policy. Finally, it provides some insight into why the optimality conditions for the single marble multi-urn search problem given in equation (8) result in a greedy policy. Equation (8) specifies that the first marble drawn is the one that maximizes the probability of immediately finding the target. It follows from Theorem 8 that subsequent draws from the same urn will continue to maximize this probability.

3.5. Urn Correlation Dynamics

In the preceding section we examined how urn probabilities and the probability of drawing a red marble from each urn evolved as a function of the state of the system. In this section we show by example how, in general, the correlations among the urns can evolve in ways that we find to be counterintuitive.

We say that two urns i and j are *positively correlated* at stage t if they have positive covariance, i.e.,

$$\mathbb{P}\left(A_i \cap A_j \mid \mathbf{x}(t)\right) > \mathbb{P}\left(A_i \mid \mathbf{x}(t)\right)\mathbb{P}\left(A_j \mid \mathbf{x}(t)\right).$$

Likewise, urns i and j are *negatively correlated* if their covariance is negative,

$$\mathbb{P}\left(A_i \cap A_j \mid \mathbf{x}(t)\right) < \mathbb{P}\left(A_i \mid \mathbf{x}(t)\right)\mathbb{P}\left(A_j \mid \mathbf{x}(t)\right).$$

In Theorem 3 we showed the somewhat intuitive result that independence among the urn probabilities is preserved through state transitions. In general, correlations can change through Bayesian updates each time a blue marble is drawn. These changes can include changes in sign.

We now provide an example scenario in which all correlations are positive in the initial conditions, but after a blue marble is drawn some correlations become negative. Suppose we have a multi-urn search problem consisting of three urns. Each urn i has $N_i = 1$ marble and initial probability $\varphi_i = \frac{1}{2}$ of containing a red marble. Furthermore,

$$\varphi_{\{1,2\}} = \varphi_{\{1,3\}} = \varphi_{\{2,3\}} = \varphi_{\{1,2,3\}} = \frac{1}{3}.$$

A quick calculation confirms that this is a valid probability model, and that a red marble exists in at least one of the three urns with probability $\frac{5}{6}$. We also verify that all correlations are positive,

$$\varphi_{\{1,2\}} > \varphi_1\varphi_2$$
$$\varphi_{\{1,3\}} > \varphi_1\varphi_3$$
$$\varphi_{\{2,3\}} > \varphi_2\varphi_3.$$

Also, we have a more general positive correlation property,

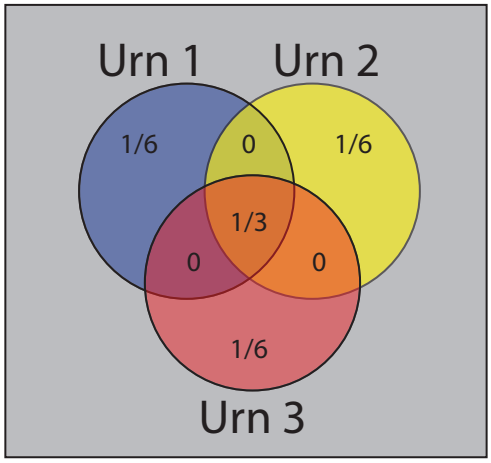

Figure 2 **Venn diagram of the probabilities in the three-urn example.**

$$\varphi_{\{1,2,3\}} > \varphi_1\varphi_2\varphi_3.$$

Figure 2 depicts this probability law in a Venn diagram. Note that there is no probability of exactly two urns containing red marbles. A red marble is present in zero, one, or three urns almost surely.

Now suppose at stage $t = 0$, a blue marble is drawn from urn 1. From Theorem 1, the stage 1 urn probabilities are

$$\mathbb{P}(A_1 \,|\, \mathbf{x}(1)) = 0$$
$$\mathbb{P}(A_2 \,|\, \mathbf{x}(1)) = \frac{1}{3}$$

$$\begin{aligned}
\mathbb{P}(A_3 \mid \mathbf{x}(1)) &= \frac{1}{3} \\
\mathbb{P}(A_1 \cap A_2 \mid \mathbf{x}(1)) &= 0 \\
\mathbb{P}(A_1 \cap A_3 \mid \mathbf{x}(1)) &= 0 \\
\mathbb{P}(A_2 \cap A_3 \mid \mathbf{x}(1)) &= 0 \\
\mathbb{P}(A_1 \cap A_2 \cap A_3 \mid \mathbf{x}(1)) &= 0.
\end{aligned}$$

By eliminating the possibility that each urn contained a red marble, the only outcomes that have positive probability in stage 1 are single-marble outcomes. The correlation between urns 2 and 3, which was positive in the initial conditions, has become negative in stage 1:

$$\mathbb{P}(A_2 \cap A_3 \mid \mathbf{x}(1)) = 0 < \frac{1}{9} = \mathbb{P}(A_2 \mid \mathbf{x}(1))\,\mathbb{P}(A_3 \mid \mathbf{x}(1)).$$

One could similarly produce examples in which correlations that were originally negative become positive after drawing blue marbles. In the single-marble and independent urn cases, for which we provided characterizations of the optimal policies in the preceding sections, correlations among the urns exhibited some stationarity with respect to stage. In general, the nature of correlations among the urns can change substantially as blue marbles are drawn and probabilities are updated. This characteristic presents a challenge to finding characterizations of the optimal policy in general.

4. Proofs of Theorems

In this section we provide the technical proof for each theorem.

4.1. Theorem 1

Proof of Theorem 1 Substituting the initial condition, $\mathbf{x}(0) = \mathbf{0}$, into the result returns the prior

$$\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(0)\right) = \varphi_U := \mathbb{P}\left(\bigcap_{i\in U} A_i\right).$$

The proof proceeds by induction. First note that in order to reach state $\mathbf{x}(t+1)$ from stage t, a blue marble must have been drawn from urn $u(t)$ from state $\mathbf{x}(t)$. We use the law of total probability to decompose this event and form a recursion:

$$\begin{aligned}
\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t+1)\right) &= \frac{\mathbb{P}\left(\left(\bigcap_{i\in U} A_i\right) \cap w(\mathbf{x}(t), u(t)) = 0 | \mathbf{x}(t)\right)}{\mathbb{P}(w(\mathbf{x}(t), u(t)) = 0)} \\
&= \frac{\left(1 - \frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}\left(\bigcap_{i\in U\cup u(t)} A_i \,\middle|\, \mathbf{x}(t)\right) + \mathbb{P}(A^c_{u(t)} \cap \left(\bigcap_{i\in U} A_i\right) | \mathbf{x}(t))}{1 - \left(\frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}(A_{u(t)} | \mathbf{x}(t))} \\
&= \frac{\mathbb{P}\left(\bigcap_{i\in U} A_i \mid \mathbf{x}(t)\right) - \left(\frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}\left(\bigcap_{i\in U\cup u(t)} A_i \,\middle|\, \mathbf{x}(t)\right)}{1 - \left(\frac{1}{N_{u(t)} - x_{u(t)}(t)}\right) \mathbb{P}(A_{u(t)} | \mathbf{x}(t))}
\end{aligned} \tag{9}$$

The result in Theorem 1 forms our induction hypothesis. We use this result to form the three probabilities in the recursion given in equation (9).

$$\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t)\right) = \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}$$

$$\mathbb{P}\left(A_{u(t)}|\mathbf{x}(t)\right) = \frac{\left(1-\frac{x_{u(t)}(t)}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\, u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}$$

$$\mathbb{P}\left(\bigcap_{i\in U\cup u(t)} A_i \,\middle|\, \mathbf{x}(t)\right)$$
$$= \begin{cases} \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}} & u(t)\in U \\ \frac{\left[\prod_{i\in (U\cup u(t))}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq (U\cup u(t))\}}(-1)^{|S|-|U|-1}\varphi_S\prod_{j\in S\setminus (U\cup u(t))}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}} & u(t)\notin U. \end{cases}$$

As we see from these probabilities, we have two cases to consider:

1. $u(t)\in U$.
2. $u(t)\notin U$.

We now look at each of these cases individually.

Case 1: $u(t)\in U$. We begin by substituting the probabilities formed using the induction hypothesis into the recursion in equation (9).

$$\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t+1)\right) = \frac{\mathbb{P}\left(\bigcap_{i\in U} A_i \,\middle|\, \mathbf{x}(t)\right) - \left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}\left(\bigcap_{i\in U\cup u(t)} A_i \,\middle|\, \mathbf{x}(t)\right)}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t))}$$
$$= \frac{\left(1-\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\left(\frac{\left(1-\frac{x_{u(t)}(t)}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\, u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\right)}$$
$$= \frac{\left(\frac{N_{u(t)}-x_{u(t)}(t)-1}{N_{u(t)}-x_{u(t)}(t)}\right)\frac{\left[\left(\frac{N_{u(t)}-x_{u(t)}(t)}{N_{u(t)}}\right)\prod_{i\in U\setminus\{u(t)\}}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}}{1-\left(\frac{1}{N_{u(t)}}\right)\left(\frac{\sum_{\{S\subseteq\mathcal{V}:\, u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\right)}$$

We proceed by separating the summations into terms corresponding to sets containing $u(t)$ and those that do not contain $u(t)$. We can then factor out the terms corresponding to urn $u(t)$ and make the following substitutions:

$$x_i(t) = \begin{cases} x_i(t+1) & i\neq u(t) \\ x_i(t+1)-1 & i=u(t). \end{cases}$$

Continuing the simplification from above,

$$
\begin{aligned}
&= \frac{\left(\frac{N_{u(t)}-x_{u(t)}(t+1)}{N_{u(t)}}\right)\left[\prod_{i\in U\setminus\{u(t)\}}\left(1-\frac{x_i(t+1)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}-\left(\frac{1}{N_{u(t)}}\right)\left(\sum_{\{S\subseteq\mathcal{V}:\,u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}\right)}\\
&= \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t+1)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}}{\sum_{S\subseteq\mathcal{V}\setminus u(t)}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}+\left(\frac{x_{u(t)}(t)}{N_{u(t)}}+\frac{1}{N_{u(t)}}\right)\left(\sum_{\{S\subseteq\mathcal{V}:\,u(t)\in S\}}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}\right)}\\
&= \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t+1)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}}{\sum_{S\subseteq\mathcal{V}\setminus u(t)}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t+1)}{N_j}+\sum_{\{S\subseteq\mathcal{V}:\,u(t)\in S\}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t+1)}{N_j}}\\
&= \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t+1)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t+1)}{N_j}}.
\end{aligned}
$$

Observe that this is the desired result for stage $t+1$. We now provide the induction step for the case in which $u(t)\notin U$.

Case 2: $u(t)\notin U$ follows a similar set of steps. We begin by substituting the probabilities formed using the induction hypothesis into the recursion in equation (9).

$$
\begin{aligned}
&\frac{\mathbb{P}\left(\bigcap_{i\in U}A_i\,\middle|\,\mathbf{x}(t)\right)-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}\left(\bigcap_{i\in U\cup u(t)}A_i\,\middle|\,\mathbf{x}(t)\right)}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t))}\\
&\quad=\left[\frac{\left(\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right)\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\right.\\
&\qquad\left.-\frac{\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\left(\prod_{i\in U\cup u(t)}\left(1-\frac{x_i(t)}{N_i}\right)\right)\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\cup u(t)\}}(-1)^{|S|-|U|-1}\varphi_S\prod_{j\in S\setminus(U\cup u(t))}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\right]\\
&\qquad\times\left[1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\left(\frac{\left(1-\frac{x_{u(t)}(t)}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\,u(t)\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{u(t)\}}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\right)\right]^{-1}
\end{aligned}
$$

The denominators in the above expression reduce in exactly the same way as in the previous case in which $u(t)\in U$. In fact, the denominators in the induction hypothesis and in equation (9) do not depend on whether $u(t)\in U$. Because we have already shown the steps for reducing this denominator to the desired form, we omit these steps and only show the induction on the numerators for this case:

$$
\begin{aligned}
&\left[\left(\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right)\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}\right.\\
&\qquad\left.-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\left(\prod_{i\in U\cup u(t)}\left(1-\frac{x_i(t)}{N_i}\right)\right)\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\cup u(t)\}}(-1)^{|S|-|U|-1}\varphi_S\prod_{j\in S\setminus(U\cup u(t))}\frac{x_j(t)}{N_j}\right]
\end{aligned}
$$

$$= \left(\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right)\left[\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}\right.$$
$$\left.+\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\left(\frac{N_{u(t)}-x_{u(t)}(t)}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\cup u(t)\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus(U\cup u(t))}\frac{x_j(t)}{N_j}\right]$$
$$= \left(\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right)\left[\sum_{\{S\subseteq\mathcal{V}\setminus\{u(t)\}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}\right.$$
$$\left.+\left(\frac{x_{u(t)}(t)}{N_{u(t)}}+\frac{1}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\cup u(t)\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus(U\cup u(t))}\frac{x_j(t)}{N_j}\right]$$

We again make the substitution:

$$x_i(t)=\begin{cases}x_i(t+1) & i\neq u(t)\\ x_i(t+1)-1 & i=u(t),\end{cases}$$

and continue from above:

$$= \left(\prod_{i\in U}\left(1-\frac{x_i(t+1)}{N_i}\right)\right)\left[\sum_{\{S\subseteq\mathcal{V}\setminus\{u(t)\}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}\right.$$
$$\left.+\left(\frac{x_{u(t)}(t+1)}{N_{u(t)}}\right)\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\cup u(t)\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus(U\cup u(t))}\frac{x_j(t+1)}{N_j}\right]$$
$$= \left(\prod_{i\in U}\left(1-\frac{x_i(t+1)}{N_i}\right)\right)\sum_{\{S\subseteq\mathcal{V}:\, S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t+1)}{N_j}.$$

This final expression is the numerator in Theorem 1 for the stage $t+1$ urn probabilities. □

4.2. Theorem 2

Before we provide a proof for Theorem 2, we state and prove the following Lemma.

Lemma 1 *Given a multi-urn search problem on a set of urns $\mathcal{V}$, suppose a policy $\mathbf{u}$ is executed to stage t irrespective of whether a red marble is found at any stage. Let B_i be the event that a red marble has been drawn from urn $i\in\mathcal{V}$ in this experiment. Then, for any subset $U\subseteq\mathcal{V}$, the probability of having drawn a red marble from each of the urns in U and none of the other urns is*

$$\mathbb{P}\left(\left(\bigcap_{i\in U}B_i\right)\cap\left(\bigcap_{j\in\mathcal{V}\setminus U}B_j^c\right)\right)=\sum_{S:S\subseteq\mathcal{V},S\supseteq U}(-1)^{|S|-|U|}\varphi_S\prod_{i\in S}\frac{x_i(t)}{N_i}\geq 0.$$

Proof of Lemma 1. This result is comes from the principle of inclusion-exclusion, and follows from Equation 1. From basic set operations and DeMorgan's Law, we can write

$$\left(\bigcap_{i\in U}B_i\right)=\left(\left(\bigcap_{i\in U}B_i\right)\cap\left(\bigcap_{j\in\mathcal{V}\setminus U}B_j^c\right)\right)\cup\left(\left(\bigcap_{i\in U}B_i\right)\cap\left(\bigcup_{j\in\mathcal{V}\setminus U}B_j\right)\right),$$

which is a union of disjoint sets. Therefore,

$$\mathbb{P}\left(\left(\bigcap_{i\in U} B_i\right)\cap\left(\bigcap_{j\in\mathcal{V}\setminus U} B_j^c\right)\right)=\mathbb{P}\left(\bigcap_{i\in U} B_i\right)-\mathbb{P}\left(\left(\bigcap_{i\in U} B_i\right)\cap\left(\bigcup_{j\in\mathcal{V}\setminus U} B_j\right)\right). \tag{10}$$

The probability a red marble is drawn from all of the urns in U in this experiment can be found using the multiplication rule:

$$\mathbb{P}\left(\bigcap_{i\in U} B_i\right)=\varphi_U\prod_{i\in U}\frac{x_i(t)}{N_i}. \tag{11}$$

Recall that φ_U is the probability of all of the urns in set U containing a red marble, and $\frac{x_i(t)}{N_i}$ is simply the fraction of marbles removed from urn i at stage t. The product in this expression implies conditional independence: given all of the urns in U contain a red marble, the probability that a red marble is drawn from each of them by stage t is the product of the individual probabilities. This conditional independence follows implicitly from our search assumptions. The order of marble draws from each urn is random, and does not depend on the order of marble draws from any other urn.

We now examine $\mathbb{P}\left(\left(\bigcap_{i\in U} B_i\right)\cap\left(\bigcup_{j\in\mathcal{V}\setminus U} B_j\right)\right)$. We first note that

$$\left(\left(\bigcap_{i\in U} B_i\right)\cap\left(\bigcup_{j\in\mathcal{V}\setminus U} B_j\right)\right)=\bigcup_{j\in\mathcal{V}\setminus U}\left(\left(\bigcap_{i\in U} B_i\right)\cap B_j\right).$$

Using the principle of inclusion exclusion, we can find the probability of this union,

$$\begin{aligned}\mathbb{P}\left(\bigcup_{j\in\mathcal{V}\setminus U}\left(\left(\bigcap_{i\in U} B_i\right)\cap B_j\right)\right)&=\sum_{j\in\mathcal{V}\setminus U}\mathbb{P}\left(\left(\bigcap_{i\in U} B_i\right)\cap B_j\right)\\&\quad-\sum_{j\in\mathcal{V}\setminus U}\left[\sum_{k\in\mathcal{V}\setminus(U\cup\{j\})}\mathbb{P}\left(\left(\bigcap_{i\in U} B_i\right)\cap B_j\cap B_k\right)\right]\\&\quad\cdots+(-1)^{|\mathcal{V}|-|U|+1}\mathbb{P}\left(\bigcap_{i\in\mathcal{V}} B_i\right).\end{aligned} \tag{12}$$

Substituting expressions 11 and 12 into equation 10 reduces to the desired result. The principle of inclusion-exclusion and the axioms of probability ensure that this quantity is nonnegative. □

We now provide the proof of Theorem 2.

Proof of Theorem 2. Suppose we are given a multi-urn search problem on a set of urns $\mathcal{V}$, with each urn $i\in\mathcal{V}$ containing N_i marbles and initial target probabilities φ_U for all $U\subseteq\mathcal{V}$. Suppose also that valid policy $\mathbf{u}=(u(0),\ldots,u(N-1))$ is optimal, where $\mathbf{u}$ is not a block policy. This implies that we can find a stage τ where

$$\begin{aligned}u(\tau)&=i\\u(\tau+1),u(\tau+2),\ldots,u(\tau+\delta-1)&\neq i\\u(\tau+\delta)&=i,\end{aligned}$$

for some urn $i \in \mathcal{V}$, where $\delta > 1$.

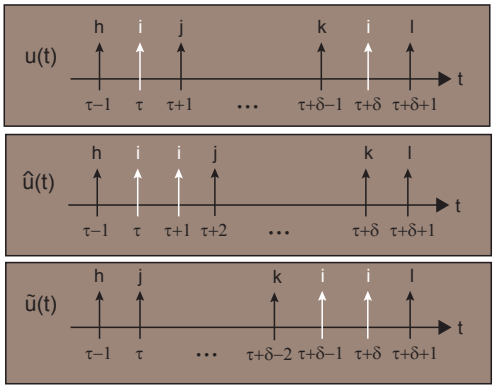

Figure 3 Comparison of policies u, û, and ũ.

We now consider two alternative policies that move the queries of urn i into "blocks". Policy $\hat{\mathbf{u}}$ executes the two queries of i in stages τ and $\tau+1$, then executes the rest of the queries in the subsequence. Policy $\tilde{\mathbf{u}}$ executes the two queries of i after executing the other queries in the subsequence. A visual comparison of these policies is provided in Figure 3. Formally,

$$\hat{\mathbf{u}} = \hat{u}(t) = \begin{cases} u(t), & t \leq \tau \text{ or } t > \tau+\delta \\ u(t-1), & t = \tau+1, \tau+2, \ldots, \tau+\delta, \end{cases}$$

$$\tilde{\mathbf{u}} = \tilde{u}(t) = \begin{cases} u(t), & t < \tau \text{ or } t \geq \tau+\delta \\ u(t+1), & t = \tau, \tau+1, \ldots, \tau+\delta-1. \end{cases}$$

Let C be the number of non-target queries (i.e., the cost) when using policy $\mathbf{u}$, $\hat{C}$ be the same quantity when using policy $\hat{\mathbf{u}}$, and $\tilde{C}$ be the same quantity when using policy $\tilde{\mathbf{u}}$. Conditioned on not having found a red marble, it follows that the state trajectories,

$$\hat{x}_j(t) = \begin{cases} x_j(t), & t \leq \tau+1 \text{ or } t > \tau+\delta \\ x_j(t-1), & t = \tau+2, \ldots, \tau+\delta,\ j \neq i \\ x_j(t-1)+1, & t = \tau+2, \ldots, \tau+\delta,\ j = i \end{cases}$$

$$\tilde{x}_j(t) = \begin{cases} x_j(t), & t <= \tau \text{ or } t \geq \tau+\delta \\ x_j(t+1), & t = \tau+1, \ldots, \tau+\delta-1,\ j \neq i \\ x_j(t+1)-1, & t = \tau+1, \ldots, \tau+\delta-1,\ j = i, \end{cases}$$

where $\hat{x}_j(t)$ and $\tilde{x}_j(t)$ are the numbers of times urn j has been queried before stage t under policies $\hat{\mathbf{u}}$ and $\tilde{\mathbf{u}}$, respectively.

Using equation (5), our assumptions imply that

$$\mathbb{E}[\hat{C}] - \mathbb{E}[C] \geq 0$$

$$\sum_{t=0}^{N-1} \sum_{S \subseteq \mathcal{V}} (-1)^{|S|} \varphi_S \prod_{j \in S} \frac{\hat{x}_j(t+1)}{N_j} - \sum_{t=0}^{N-1} \sum_{S \subseteq \mathcal{V}} (-1)^{|S|} \varphi_S \prod_{j \in S} \frac{x_j(t+1)}{N_j} \geq 0$$

$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left(\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\hat{x}_j(t)}{N_j}-\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\hat{x}_j(t)}{N_j}-\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left(\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}-\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\left(\frac{x_i(t)+1}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}-\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\left(\frac{x_i(t)}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\right)\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left((-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+1)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+\delta)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\frac{x_i(\tau+\delta)}{N_i}\right)\left((-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}\right)$$
$$+\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left((-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+1)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\frac{x_i(\tau+\delta)}{N_i}\right)\left((-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}\right)$$
$$-\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}\sum_{t=\tau+2}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\leq\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}.$$

Optimality of $\mathbf{u}$ likewise implies

$$\mathbb{E}[\tilde{C}]-\mathbb{E}[C]\geq 0$$
$$\sum_{t=0}^{N-1}\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\tilde{x}_j(t+1)}{N_j}-\sum_{t=0}^{N-1}\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t+1)}{N_j}\geq 0$$
$$\sum_{S\subseteq\mathcal{V}}\left(\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\tilde{x}_j(t)}{N_j}-\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left(\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\tilde{x}_j(t)}{N_j}-\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{\tilde{x}_j(t)}{N_j}-\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left(\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}-\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\sum_{t=\tau+2}^{\tau+\delta}(-1)^{|S|}\varphi_S\left(\frac{x_i(t)-1}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}-\sum_{t=\tau+1}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\left(\frac{x_i(t)}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\right)\geq 0$$
$$\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left((-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+1)}{N_j}\right)$$
$$+\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\varphi_S\left(\frac{x_i(\tau+\delta)-1}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\left(\frac{x_i(\tau+1)}{N_i}\right)\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}$$

$$
\begin{aligned}
&-\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}\sum_{t=\tau+2}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\geq 0\\
&\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left((-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+1)}{N_j}\right)\\
&+\sum_{S\subseteq\mathcal{V}:i\in S}\frac{x_i(\tau+\delta)}{N_i}\left((-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}\right)\\
&-\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}\sum_{t=\tau+2}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\geq\left(\frac{1}{N_i}\right)\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}.
\end{aligned}
$$

Now let

$$
\begin{aligned}
\alpha=&\sum_{S\subseteq\mathcal{V}\setminus\{i\}}\left((-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(\tau+1)}{N_j}\right)\\
&+\sum_{S\subseteq\mathcal{V}:i\in S}\left(\frac{x_i(\tau+\delta)}{N_i}\right)\left((-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}-(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}\right)\\
&-\frac{1}{N_i}\sum_{S\subseteq\mathcal{V}:i\in S}\sum_{t=\tau+2}^{\tau+\delta-1}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j},
\end{aligned}
$$

which appears in both of the expected cost inequalities we have derived. We have established that

$$
\begin{aligned}
\alpha&\geq\left(\frac{1}{N_i}\right)\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+\delta)}{N_j}\\
\alpha&\leq\left(\frac{1}{N_i}\right)\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(\tau+1)}{N_j}
\end{aligned}
$$

We now show that for any $i\in\mathcal{V}$, for any $u(t)\in\mathcal{V}\setminus\{i\}$,

$$
h(t)=\sum_{S\subseteq\mathcal{V}:\,i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j} \tag{13}
$$

is a nondecreasing function that is strictly increasing when $\varphi_{\{i,u(t)\}}>0$.

Observe that

$$
\begin{aligned}
h(t+1)&=\sum_{S\subseteq\mathcal{V}:\,i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t+1)}{N_j}\\
&=\sum_{S\subseteq\mathcal{V}\setminus\{u(t)\}:\,i\in S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}\\
&\quad+\left(\frac{x_{u(t)}(t)+1}{N_{u(t)}}\right)\sum_{S\subseteq\mathcal{V}:\,\{i,u(t)\}\subseteq S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i,u(t)\}}\frac{x_j(t)}{N_j}\\
&=h(t)+\left(\frac{1}{N_{u(t)}}\right)\sum_{S\subseteq\mathcal{V}:\,\{i,u(t)\}\subseteq S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i,u(t)\}}\frac{x_j(t)}{N_j}.
\end{aligned}
$$

From Lemma 1,

$$\left(\frac{x_i(t)x_{u(t)}(t)}{N_iN_{u(t)}}\right)\sum_{S\subseteq\mathcal{V}:\,\{i,u(t)\}\subseteq S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i,u(t)\}}\frac{x_j(t)}{N_j}\geq 0$$
$$\Rightarrow\left(\frac{1}{N_{u(t)}}\right)\sum_{S\subseteq\mathcal{V}:\,\{i,u(t)\}\subseteq S}(-1)^{|S|}\varphi_S\prod_{j\in S\setminus\{i,u(t)\}}\frac{x_j(t)}{N_j}\geq 0.$$

This final inequality implies $h(t+1)\geq h(t)$. Because urn i is not queried in stages $t=\tau+1,\ldots,\tau+\delta$, $h(t)$ is nondecreasing over these stages. Therefore we can write

$$\alpha\geq h(\tau+\delta)\geq h(\tau+1)\geq\alpha. \tag{14}$$

This expression can only be satisfied by equality, which means that for any optimal policy that is not a block policy, we can maintain optimality while successively permuting the policy so that the urn queries are arranged into blocks. □

4.3. Theorem 3

Proof of Theorem 3. This result comes from substituting the appropriate products into the result from Theorem 1. For any urn $i\in\mathcal{V}$,

$$\begin{aligned}
\mathbb{P}(A_i|\mathbf{x}(t)) &= \frac{\left(1-\frac{x_i(t)}{N_i}\right)\sum_{\{S\subseteq\mathcal{V}:\,i\in S\}}(-1)^{|S|-1}\varphi_S\prod_{j\in S\setminus\{i\}}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\\
&= \frac{\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\sum_{\{S\subseteq\mathcal{V}\setminus\{i\}\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}\\
&= \frac{\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\sum_{\{S\subseteq\mathcal{V}\setminus\{i\}\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}\setminus\{i\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}+\sum_{S\subseteq\mathcal{V}:i\in S}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}\\
&= \frac{\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\sum_{\{S\subseteq\mathcal{V}\setminus\{i\}\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}\setminus\{i\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}-\varphi_i\frac{x_i(t)}{N_i}\sum_{S\subseteq\mathcal{V}\setminus\{i\}}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}\\
&= \frac{\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)}{1-\varphi_i\left(\frac{x_i(t)}{N_i}\right)}.
\end{aligned}$$

In a similar manner we use Theorem 1 to find the urn probability for subset $U\subseteq\mathcal{V}$ at stage t,

$$\begin{aligned}
\mathbb{P}\left(\bigcap_{i\in U}A_i\right) &= \frac{\left[\prod_{i\in U}\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{\{S\subseteq\mathcal{V}:\,S\supseteq U\}}(-1)^{|S|-|U|}\varphi_S\prod_{j\in S\setminus U}\frac{x_j(t)}{N_j}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S\prod_{j\in S}\frac{x_j(t)}{N_j}}\\
&= \frac{\left[\prod_{i\in U}\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{S\subseteq\mathcal{V}\setminus U}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}{\sum_{T\subseteq U}\sum_{S\subseteq\mathcal{V}\setminus U}(-1)^{|S|+|T|}\prod_{j\in S\cup T}\varphi_j\frac{x_j(t)}{N_j}}
\end{aligned}$$

$$= \frac{\left[\prod_{i\in U}\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\right]\sum_{S\subseteq\mathcal{V}\setminus U}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}{\sum_{T\subseteq U}(-1)^{|T|}\prod_{k\in T}\varphi_k\frac{x_k(t)}{N_k}\sum_{S\subseteq\mathcal{V}\setminus U}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}}$$
$$= \frac{\left[\prod_{i\in U}\varphi_i\left(1-\frac{x_i(t)}{N_i}\right)\right]}{\sum_{S\subseteq U}(-1)^{|S|}\prod_{j\in S}\varphi_j\frac{x_j(t)}{N_j}} = \prod_{i\in U}\mathbb{P}(A_i).$$

In the final equality, we have used the property that for any $\beta_1,\ldots,\beta_M$,

$$\prod_{i=1}^{M}(1-\beta_i) = \sum_{S\subseteq[M]}\prod_{j\in S}(-\beta_j).$$

We can verify this property by induction. Define $\prod_{j\in\emptyset}(-\beta_j)=1$. Now observe

$$\begin{aligned}\prod_{i=1}^{M+1}(1-\beta_i) &= (1-\beta_{M+1})\prod_{i=1}^{M}(1-\beta_i)\\ &= (1-\beta_{M+1})\sum_{S\subseteq[M]}\prod_{j\in S}(-\beta_j)\\ &= \sum_{S\subseteq[M]}\prod_{j\in S}(-\beta_j) - \beta_{M+1}\sum_{S\subseteq[M]}\prod_{j\in S}(-\beta_j)\\ &= \sum_{S\subseteq[M+1]}\prod_{j\in S}(-\beta_j).\end{aligned}$$

Setting $\beta_i = \varphi_i\left(\frac{x_i(t)}{N_i}\right)$ achieves the desired result. □

4.4. Theorem 4

Proof of Theorem 4. First we prove that the condition in equation 7 implies optimality by contrapositive. Let $\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$ be a block policy that does not satisfy this condition, and let i be an index for which

$$N_{v^i}\left(\frac{2-\varphi_{v^i}}{\varphi_{v^i}}\right) > N_{v^{i+1}}\left(\frac{2-\varphi_{v^{i+1}}}{\varphi_{v^{i+1}}}\right).$$

Also, we define the first stage in which urn v^i is queried in this policy as $\tau = \sum_{j=0}^{i-1}N_j$.

We now construct an alternative block policy $\tilde{\mathbf{u}}_B$, so that

$$\tilde{v}^j = \begin{cases} v^j & j\notin\{i,i+1\}\\ v^{i+1} & j=i\\ v^i & j=i+1.\end{cases}$$

Let $\mathbb{E}[C] = \sum_{k=0}^{N-1}\prod_{t=0}^{k}\mathbb{P}(w(\mathbf{x}(t),u(t))=0)$ be the expected cost of policy $\mathbf{u}_B$ and $\mathbb{E}[\tilde{C}] = \sum_{k=0}^{N-1}\prod_{t=0}^{k}\mathbb{P}(w(\tilde{\mathbf{x}}(t),\tilde{u}(t))=0)$ be the expected cost of policy $\tilde{\mathbf{u}}_B$. For brevity, we also define

$\gamma = \prod_{t=0}^{\tau-1} P(w(\mathbf{x}(t), u(t)) = 0) = \prod_{j=1}^{i-1}(1-\varphi_{v^j}) > 0$, according to Corollary 4. Now consider the difference in expected cost,

$$\begin{aligned}
\mathbb{E}[C] - \mathbb{E}[\tilde{C}] &= \sum_{k=0}^{N-1}\prod_{t=0}^{k}\mathbb{P}(w(\mathbf{x}(t),u(t))=0) - \sum_{k=0}^{N-1}\prod_{t=0}^{k}\mathbb{P}(w(\tilde{\mathbf{x}}(t),\tilde{u}(t))=0) \\
&= \sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1}\prod_{t=0}^{k}\mathbb{P}(w(\mathbf{x}(t),u(t))=0) \\
&\quad - \sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1}\prod_{t=0}^{k}\mathbb{P}(w(\tilde{\mathbf{x}}(t),\tilde{u}(t))=0) \\
&\overset{(a)}{=} \gamma\left(\sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1}\prod_{t=\tau}^{k}\mathbb{P}(w(\mathbf{x}(t),u(t))=0)\right) \\
&\quad - \gamma\left(\sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1}\prod_{t=\tau}^{k}\mathbb{P}(w(\tilde{\mathbf{x}}(t),\tilde{u}(t))=0)\right) \\
&\overset{(b)}{=} \gamma\left(\left(N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2}\right) + (1-\varphi_{v^i})\left(N_{v^{i+1}} - \frac{(N_{v^{i+1}}+1)\varphi_{v^{i+1}}}{2}\right)\right) \\
&\quad - \gamma\left(\left(N_{v^{i+1}} - \frac{(N_{v^{i+1}}+1)\varphi_{v^{i+1}}}{2}\right) + (1-\varphi_{v^{i+1}})\left(N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2}\right)\right) \\
&= \gamma\left(\varphi_{v^{i+1}}\left(N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2}\right) - \varphi_{v^i}\left(N_{v^{i+1}} - \frac{(N_{v^{i+1}}+1)\varphi_{v^{i+1}}}{2}\right)\right) \\
&= \gamma\varphi_{v^i}\varphi_{v^{i+1}}\left(\left(\frac{2N_{v^i} - N_{v^i}\varphi_{v^i} - \varphi_{v^i}}{2\varphi_{v^i}}\right) - \left(\frac{2N_{v^{i+1}} - N_{v^{i+1}}\varphi_{v^{i+1}} - \varphi_{v^i}}{2\varphi_{v^{i+1}}}\right)\right) \\
&= \frac{\gamma\varphi_{v^i}\varphi_{v^{i+1}}}{2}\left(N_{v^i}\left(\frac{2-\varphi_{v^i}}{\varphi_{v^i}}\right) - N_{v^{i+1}}\left(\frac{2-\varphi_{v^{i+1}}}{\varphi_{v^{i+1}}}\right)\right) > 0. \qquad (15)
\end{aligned}$$

Steps (a) and (b) follow immediately from Corollary 4. The difference in expected cost, $\mathbb{E}[C]-\mathbb{E}[\tilde{C}]$, is strictly positive so that block policy $\mathbf{u}_B$ cannot be optimal. Therefore, the condition given in Theorem 4 is necessary for optimality.

Now we show that the same condition is sufficient for optimality, i.e., if a block policy satisfies equation (7) in Theorem 4, then it must be an optimal policy. In order to form a contradiction, suppose now that $\mathbf{u}_B$ is a block policy that satisfies the condition but is not optimal. Let $\mathbf{u}_B^\star$ be the optimal block policy, which Theorem 2 guarantees to exist.

We know from the above argument that $\mathbf{u}_B^\star$ also must satisfy the condition, which implies that policies $\mathbf{u}_B$ and $\mathbf{u}_B^\star$ can only differ by permuting subsequences $v^i, v^{i+1}, \dots v^{i+\delta}$ for which

$$N_{v^i}\left(\frac{2-\varphi_{v^i}}{\varphi_{v^i}}\right) = N_{v^{i+1}}\left(\frac{2-\varphi_{v^{i+1}}}{\varphi_{v^{i+1}}}\right) = \cdots = N_{v^{i+\delta}}\left(\frac{2-\varphi_{v^{i+\delta}}}{\varphi_{v^{i+\delta}}}\right).$$

The optimal block policy $\mathbf{u}_B^\star$ therefore can be constructed by executing a finite number of sequential pairwise exchanges in the urn ordering in block policy $\mathbf{u}_B$, each satisfying the condition with

equality. However, it follows from the inequality in equation (15) that any such permutation results in the same expected policy cost. We can conclude that the expected costs of the two policies are equal, establishing the contradiction and showing $\mathbf{u}_B$ to be an optimal policy. □

4.5. Theorem 6

Proof of Theorem 6. The proof is similar to that of Theorem 4. First we prove that the condition in equation (8) implies optimality by contrapositive. Let $\mathbf{u}_B = (v^1, v^2, \ldots, v^{|\mathcal{V}|})$ be a block policy that does not satisfy equation (8), and let i be an index for which

$$\frac{\varphi_{v^i}}{N_{v^i}} < \frac{\varphi_{v^{i+1}}}{N_{v^{i+1}}}.$$

Also, we define the first stage in which urn v^i is queried in this policy as $\tau = \sum_{j=1}^{i-1} N_j$.

We now construct an alternative block policy $\tilde{\mathbf{u}}_B$, so that

$$\tilde{v}^j = \begin{cases} v^j & j \notin \{i, i+1\} \\ v^{i+1} & j = i \\ v^i & j = i+1. \end{cases}$$

Let $\mathbb{E}[C] = \sum_{k=0}^{N-1} \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0)$ be the expected cost of policy $\mathbf{u}_B$ and $\mathbb{E}[\tilde{C}] = \sum_{k=0}^{N-1} \prod_{t=0}^{k} \mathbb{P}(w(\tilde{\mathbf{x}}(t), \tilde{u}(t)) = 0)$ be the expected cost of policy $\tilde{\mathbf{u}}_B$. Now consider the difference in expected cost,

$$\begin{aligned}
\mathbb{E}[C] - \mathbb{E}[\tilde{C}] &= \sum_{k=0}^{N-1} \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) - \sum_{k=0}^{N-1} \prod_{t=0}^{k} \mathbb{P}(w(\tilde{\mathbf{x}}(t), \tilde{u}(t)) = 0) \\
&= \sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1} \prod_{t=0}^{k} \mathbb{P}(w(\mathbf{x}(t), u(t)) = 0) \\
&\quad - \sum_{k=\tau}^{\tau+N_{v^i}+N_{v^{i+1}}-1} \prod_{t=0}^{k} \mathbb{P}(w(\tilde{\mathbf{x}}(t), \tilde{u}(t)) = 0) \\
&\overset{(c)}{=} \left(N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2} - N_{v^i} \sum_{j=1}^{i-1} \varphi_{v^j} + N_{v^{i+1}} - \frac{(N_{v^{i+1}}+1)\varphi_{v^{i+1}}}{2} - N_{v^{i+1}} \sum_{j=1}^{i} \varphi_{v^j} \right) \\
&\quad - \left(N_{v^{i+1}} - \frac{(N_{v^{i+1}}+1)\varphi_{v^{i+1}}}{2} - N_{v^{i+1}} \sum_{j=1}^{i-1} \varphi_{v^j} \right. \\
&\qquad \left. + N_{v^i} - \frac{(N_{v^i}+1)\varphi_{v^i}}{2} - N_{v^i} \sum_{j=1}^{i-1} \varphi_{v^j} - N_{v^i}\varphi_{v^{i+1}} \right) \\
&= N_{v^i}\varphi_{v^{i+1}} - N_{v^{i+1}}\varphi_{v^i} \\
&= N_{v^i} N_{v^{i+1}} \left(\frac{\varphi_{v^{i+1}}}{N_{v^{i+1}}} - \frac{\varphi_{v^i}}{N_{v^i}} \right) > 0.
\end{aligned}$$

Step (c) follows from substituting the results in Corollary 6. The difference in expected cost, $\mathbb{E}[C] - \mathbb{E}[\tilde{C}]$, is strictly positive so that block policy $\mathbf{u}_B$ cannot be optimal. Therefore, the condition given in equation (7) is necessary for optimality.

Now we show that the same condition is sufficient for optimality, i.e., if a block policy satisfies equation (8), then it must be an optimal policy. In order to form a contradiction, suppose now that $\mathbf{u}_B$ is a block policy that satisfies the condition but is not optimal. Let $\mathbf{u}_B^\star$ be the optimal block policy, which Theorem 2 guarantees to exist.

We know from the above argument that $\mathbf{u}_B^\star$ also must satisfy the condition, which implies that policies $\mathbf{u}_B$ and $\mathbf{u}_B^\star$ can only differ by permuting subsequences $v^i, v^{i+1}, \dots v^{i+\delta}$ for which

$$\frac{\varphi_{v^i}}{N_{v^i}} = \frac{\varphi_{v^{i+1}}}{N_{v^{i+1}}} = \cdots = \frac{\varphi_{v^{i+\delta}}}{N_{v^{i+\delta}}}.$$

The optimal block policy $\mathbf{u}_B^\star$ therefore can be constructed by executing a finite number of sequential pairwise exchanges in the urn ordering in block policy $\mathbf{u}_B$, in which the two urns in each exchange satisfy the equation (8) with equality. However, it follows from our previous argument that any such permutation does not affect expected policy cost. We can conclude that the expected costs of the two policies are equal, establishing the contradiction and showing $\mathbf{u}_B$ to be an optimal policy.

4.6. Theorem 7

Proof of Theorem 7. The first inequality in Theorem 7 follows from Equation 9 in Appendix 4.1. Given $u(t) \in U$,

$$\begin{aligned}\mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t+1)\right) &= \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right)\left(\frac{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t))}\right)\\ &\leq \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right).\end{aligned}$$

Note that if $\mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right) = 0$, then $\mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t+1)\right) = 0$. Equality is likewise preserved when $\mathbb{P}\left(A_{u(t)}|\mathbf{x}(t)\right) = 1$. (We intentionally omit the case when no red marble is found after drawing all of the marbles in an urn i for which $\varphi_i = 1$.) Assuming $0 < \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right)$ *and* $\mathbb{P}\left(A_{u(t)}|\mathbf{x}(t)\right) < 1$, then the inequality becomes strict:

$$\begin{aligned}0 < \left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t)) &< \left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right) \leq 1\\ 1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t)) &> 1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\\ 1 &> \frac{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t))}\\ \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right) &> \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t)\right)\left(\frac{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)}{1-\left(\frac{1}{N_{u(t)}-x_{u(t)}(t)}\right)\mathbb{P}(A_{u(t)}|\mathbf{x}(t))}\right)\\ &= \mathbb{P}\left(\bigcap_{i\in U} A_i|\mathbf{x}(t+1)\right)\end{aligned}$$

To prove the second inequality in Theorem 7, first note that for any real numbers α, β, such that $\alpha > 0$, $\alpha + \beta > 0$, and $|\beta| > 0$,

$$(\alpha+\beta)^2 = \alpha^2 + 2\alpha\beta + \beta^2 > \alpha^2 + 2\alpha\beta = \alpha(\alpha+2\beta)$$
$$\left(\frac{\alpha+\beta}{\alpha}\right) > \left(\frac{\alpha+2\beta}{\alpha+\beta}\right)$$

Now let

$$\alpha = \sum_{S\subseteq\mathcal{V}} (-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t)}{N_i}$$
$$\beta = \left(\frac{1}{N_{u(t)}}\right) \sum_{S\subseteq\mathcal{V}:u(t)\in S} (-1)^{|S|}\varphi_S \prod_{i\in S\setminus\{u(t)\}} \frac{x_i(t)}{N_i}.$$

Observe that

$$\alpha+\beta = \sum_{S\subseteq\mathcal{V}} (-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t+1)}{N_i}$$
$$\alpha+2\beta = \sum_{S\subseteq\mathcal{V}} (-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t+2)}{N_i},$$

assuming urn $u(t)$ is queried again in stage $t+1$. From Lemma 1 and Corollary 2, we can conclude that $\alpha > 0$ and $\alpha + \beta > 0$ as long as there is a positive probability of reaching stage t without finding a red marble. Therefore,

$$\left(\frac{\alpha+\beta}{\alpha}\right) \geq \left(\frac{\alpha+2\beta}{\alpha+\beta}\right)$$
$$\frac{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t+1)}{N_i}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t)}{N_i}} \geq \frac{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t+2)}{N_i}}{\sum_{S\subseteq\mathcal{V}}(-1)^{|S|}\varphi_S \prod_{i\in S} \frac{x_i(t+1)}{N_i}}$$
$$\mathbb{P}\left(w(\mathbf{x}(t), u(t)) = 0\right) \geq \mathbb{P}\left(w(\mathbf{x}(t+1), u(t)) = 0\right)$$
$$\mathbb{P}\left(w(\mathbf{x}(t), u(t)) = 1\right) \leq \mathbb{P}\left(w(\mathbf{x}(t+1), u(t)) = 1\right).$$

It follows from Lemma 1 that $\beta = 0$ only when there is no probability of drawing a red marble from urn $u(t)$ in stage t. Therefore, the inequality is strict whenever $\mathbb{P}(w(\mathbf{x}(t), u(t)) = 1) > 0$.

4.7. Theorem 8

Proof of Theorem 8. We assume that all probabilities are positive. Note that the result in Theorem 8 can be restated

$$\frac{\mathbb{P}\left(w(\mathbf{x}(t), i) = 1\right)}{\mathbb{P}\left(w(\mathbf{x}(t), u(t)) = 1\right)} \geq \frac{\mathbb{P}\left(w(\mathbf{x}(t+1), i) = 1\right)}{\mathbb{P}\left(w(\mathbf{x}(t+1), u(t)) = 1\right)}. \tag{16}$$

Now let

$$\gamma_i = \sum_{S\subseteq\mathcal{V}\setminus\{u(t)\}:\, i\in S} (-1)^{|S|}\varphi_S \prod_{i\in S\setminus\{i\}} \frac{x_i(t)}{N_i}$$
$$\gamma_{u(t)} = \sum_{S\subseteq\mathcal{V}\setminus\{i\}:\, u(t)\in S} (-1)^{|S|}\varphi_S \prod_{i\in S\setminus\{u(t)\}} \frac{x_i(t)}{N_i}$$
$$\gamma_{i,u(t)} = \sum_{S\subseteq\mathcal{V}:\, \{i,u(t)\}\subseteq S} (-1)^{|S|}\varphi_S \prod_{i\in S\setminus\{i,u(t)\}} \frac{x_i(t)}{N_i}$$

Substituting the probability distribution from Corollary 1 into equation (16), we have

$$\frac{-\left(1/N_i\right)\gamma_i - \left(\frac{x_{u(t)}(t)}{N_i N_{u(t)}}\right)\gamma_{i,u(t)}}{-\left(1/N_{u(t)}\right)\gamma_{u(t)} - \left(\frac{x_i(t)}{N_i N_{u(t)}}\right)\gamma_{i,u(t)}} \geq \frac{-\left(1/N_i\right)\gamma_i - \left(\frac{x_{u(t)}(t)+1}{N_i N_{u(t)}}\right)\gamma_{i,u(t)}}{-\left(1/N_{u(t)}\right)\gamma_{u(t)} - \left(\frac{x_i(t)}{N_i N_{u(t)}}\right)\gamma_{i,u(t)}}$$
$$-\left(\frac{x_{u(t)}(t)}{N_i N_{u(t)}}\right)\gamma_{i,u(t)} \geq -\left(\frac{x_{u(t)}(t)+1}{N_i N_{u(t)}}\right)\gamma_{i,u(t)}$$
$$0 \leq \gamma_{i,u(t)},$$

which is true from Lemma 1. □

5. Conclusion

We have presented a multi-urn search problem as a model for searching for a specific vertex in a network. Using this model, we have shown that there is always an optimal block policy in searches that meet the multi-urn search problem assumptions, irrespective of correlations in the probability model. We have also provided necessary and sufficient conditions for block policy optimality in two specific cases: independent urns and the single red marble scenario. Finally, we gave a few properties of the dynamics of the multi-urn search problem and commented on the challenges of finding more general optimality conditions.

Future Work

There are additional generalizations and extensions that we have not considered here, but which might also have interesting applications in modeling search. One such generalization is removing the constraint that an urn can have at most one red marble. This generalization might be an appropriate model for a network vertex search problem in which the network structure allowed for multiple edges between a pair of nodes. Allowing for multiple red marbles in a single urn substantially changes the dynamics of the multi-urn search problem.

Another area of further inquiry could involve examination of the performance of different policies under various urn probability models. We have shown, for example, that a purely greedy policy is not optimal in the case of independent urns. However, the counter-example suggests that there could be some lower bounds on greedy policy performance, which might depend on the total

number of urns and total number of marbles. At the very least, there appears to be limits on how suboptimal we can make a greedy policy when constructing a two-urn data set. Development in this direction could build on the results presented in [5].

We have assumed throughout our analysis that the target of the search would be easily identifiable to the searcher. In our urn model this assumption translated to clear color distinction, so we assume we know immediately whether a drawn marble is red or blue. However, we could relax this assumption in several ways. We could, for example, characterize each marble with a feature set and develop a probability model that gives us the probability that a marble is a red marble, given its set of features. Depending on the context of the search, this type of incomplete information model could evolve into a stopping problem in which the objective is to determine the best time to stop drawing new marbles in search of a red one. Multi-arm bandit models and sequential mutual information maximization [5] might offer useful approaches in this scenario.

A slight variation from this imperfect information approach would be to represent some marbles as being more "red" than others, according to some probability model. In this model, the searcher receives a reward at the end of the search that is a function of the most "red" marble found, but still has to pay a fixed cost for each draw. Like the imperfect information model, this formulation would ultimately be a stopping problem, balancing the current reward attained against the likelihood of attaining a higher reward by drawing more marbles.

We have also assumed uniform urn costs. It is plausible, however, that in some cases it might cost more to draw marbles from some urns than from others. Or, there could be a one-time access cost in order to gain the ability to draw marbles from the urn, which could represent law enforcement having to get a warrant to obtain internet or phone records for an individual.

Our assumption that marbles are drawn in a random sequence from each urn might not be valid in some settings. The Google search engine, for example, returns the most relevant results first, so that if a user does not find what he is looking for in the first few pages of results it might make sense to try a different query rather than look through the remainder of the pages. If we changed our model so that red marbles were more likely to be drawn first in each urn, then the problem would involve deciding when to stop querying an urn and switch to one that might be more promising.

References

[1] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

[2] Pavel Berkhin. A survey on pagerank computing. *Internet Mathematics*, 2(1):73–120, 2005. doi: 10.1080/15427951.2005.10129098. URL `http://dx.doi.org/10.1080/15427951.2005.10129098`.

[3] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000. ISBN 1886529094.

[4] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *CoRR*, abs/1204.5721, 2012. URL http://arxiv.org/abs/1204.5721.

[5] Yuxin Chen, S Hamed Hassani, Amin Karbasi, and Andreas Krause. Sequential information maximization: When is greedy near-optimal? In *Proc. International Conference on Learning Theory (COLT)*, 2015.

[6] Fan Chung, Shirin Handjani, and Doug Jungreis. Generalizations of polya's urn problem. *Annals of Combinatorics*, 7(2):141–153, 2003. ISSN 0219-3094. doi: 10.1007/s00026-003-0178-y. URL http://dx.doi.org/10.1007/s00026-003-0178-y.

[7] Doug Downey, Oren Etzioni, and Stephen Soderland. A probabilistic model of redundancy in information extraction. Technical report, DTIC Document, 2006.

[8] Esther Frostig and Gideon Weiss. Four proofs of gittins' multiarmed bandit theorem. *Annals of Operations Research*, 241(1):127–165, 2016. ISSN 1572-9338. doi: 10.1007/s10479-013-1523-0. URL http://dx.doi.org/10.1007/s10479-013-1523-0.

[9] JC Gittins and DM Jones. A dynamic allocation index for new-product chemical research. *report), CUED/A-Mat Stud/TR13, Department of Engineering, Cambridge University*, 1974.

[10] John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 148–177, 1979.

[11] Jessica Guynn and Elizabeth Weis. Twitter suspends 125,000 ISIL-related accounts. *USA Today*, February 6, 2016. URL http://www.usatoday.com/story/tech/news/2016/02/05/twitter-suspends-125000-isil-related-accounts/79889892/. http://www.usatoday.com/story/tech/news/2016/02/05/twitter-suspends-125000-isil-related-accounts/79889892/; Accessed April 12, 2016.

[12] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.

[13] Retsef Levi, Thomas Magnanti, and Yaron Shaposhnik. Scheduling with testing. Technical report, Working paper, 2016.

[14] H. Mahmoud. *Polya Urn Models*. Chapman & Hall/CRC Texts in Statistical Science. CRC Press, 2008. ISBN 9781420059847. URL https://books.google.com/books?id=7Bizo28c2LQC.

[15] L. J. Wei. The generalized polya's urn design for sequential medical trials. *The Annals of Statistics*, 7(2):291–296, 1979. ISSN 00905364. URL http://www.jstor.org/stable/2958811.